

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
A.I. LABORATORY

July 1972

Artificial Intelligence
Memo No. 264

DEVELOPING A MUSICAL EAR: A NEW EXPERIMENT

Jeanne Bamberger*

Work reported herein was conducted at the Artificial Intelligence Laboratory, a Massachusetts Institute of Technology research program supported in part by the Advanced Research Projects Agency of the Department of Defense and monitored by the Office of Naval Research under Contract Number N00014-70-A-0362-0002.

* Research Associate at the Artificial Intelligence Laboratory

I would like to report on some ideas we have been developing at M.I.T. for self-paced, independent music study. The aim of our approach is to nurture in students that enigmatic quality called, "musical"--be it a "musical ear" or an individual's capacity to give a "musical performance". While all of us cherish these qualities, rarely do we come to grips with them directly in teaching. More often we rely on our magical or mystical faith in the inspiration of music, itself, and its great artists, to do the teaching. And for some (maybe ultimately all) this is the best course. But what about the others to whom we teach only the techniques of playing instruments or some "facts" about music--its forms, its history and its apparent elements? How often do we have or take the time to examine the assumptions underlying these "facts" we teach, or to question the relation between what we teach and what we do as musicians?

I would like to suggest that the quality we call musical is partly a function of that individual's capacity to think musically. By this I mean his ability to project in performance and to be moved by, the groupings, transformations, and varying functions of musical events in a given work. And especially his ability to "perceive" the subtle and dynamic relation of detail as it becomes part of larger design--part as related to whole.

With this goal of musical thinking in mind, then, we are developing a learning environment in which students can handle, manipulate and transform musical ideas that intertwine directly with the creation of their own musical facts and artifacts. We ask students to confront real and sometimes problematic musical situations of their own and others' making; to explore and question their own results.

Our work thus far has taught us, however, that we are working with remarkably complex phenomena and equally complex human responses. And yet, paradoxically, the performance of a piece of music or a listener's experience seems to be, as it occurs, immediate, direct and all-at-once--whole. How this happens, what one needs to learn to make it so, we still know very little about. But it seems clear that such learning is also complex and multi-leveled. It is not surprising, then, that the means we find for nurturing these abilities may be far from complete at any particular step along the way. But this should not defeat the search for an approach and materials that are direct and intuitive in their effect. Nor should we capitulate to teaching "rituals" or "canned routines" which may be easier but which can also distort both the subject matter and our ultimate goals.

The approach we are developing rests on two fundamental notions derived from our own work and from the work of psychologists in the field of cognitive development.*

1) Perception and cognition are inextricably intertwined--that is, what is casually termed "perception" is not a passive taking in of phenomena but an

*See, for example, Piaget, J., Psychology of Intelligence, Littlefield, Adams, & Co., Totowa, N.J., 1966; and Inhelder, B. & Piaget, J., The Early Growth of Logic in The Child, W.W. Norton & Co., New York, 1969; Bruner, Goodnow, Austin, A STUDY OF THINKING, SCIENCE EDITIONS, INC., New York, 1962.

active organizing process in which the listener discovers or constructs coherence by spontaneously and/or deliberately processing or coding the phenomena before him. Thus, the perception of music will vary as a listener's available "categories" lead him to seize on different aspects of what comes at him. From this derives the assumption that such responses as liking or not liking a particular composition, the decisions of the composer or performer as well as all sorts of affective response cannot be separated from what the individual is actually able to "hear"--that is, his particular capacity for processing the events and their relationships within a given composition.

2) Learning behavior in music follows similar developmental patterns as that in learning generally. Thus, studies of cognitive growth and its relation to perception should be relevant to teaching and learning music too. Therefore, you will find in the means we are developing for teaching music:

- (a) An emphasis on concrete handling and manipulating of the various dimensions of music--pitch, time, sonority, structure--and their interrelationships. And because "concrete" in relation to music must refer (perhaps paradoxically) to concrete sound and time, manipulating and handling must be primarily by ear rather than by eye.
- (b) An emphasis on generative concepts--that is, an effort to give students initial "primitives" that are as extendable as possible. This is much like teaching a child to deal with numbers so that he is not limited initially to $1+2$ but can quickly extend that to $10+20$ or $100+200$.

- (c) An emphasis on music as a dynamic process, particularly as one thinks about or learns the "meaning" of pitch relations (intervals, chords), rhythm (beat, durations, patterns), or structure (phrase, forms). This in contrast to an emphasis on absolute definition as a goal of learning—for example, naming or identification of discrete or isolated bits outside of their role within the dynamic process of a specific musical moment.
- (d) As a corollary to (c), a focus on contextual meanings: Learning proceeds from the student's intuitive ability to perceive contextual relations and from his corresponding affective response. The student should be provided with an environment and analytic tools whereby he can make conscious, and thereby develop this initial intuitive sense of structure. He should learn to influence, not inhibit his own intuition.

While we are currently working with computer-related gadgets, the substance of the approach is certainly as appropriate to more conventional and more readily available means. The use of a computer controlled music player does have, however, certain nice advantages. The student can, by listening, handle and create whole musical structures without first developing the kind of facility with an instrument that would usually be necessary for such activities. This means that a student can immediately consider the results of his musical thought and aural imagination. He does not have to wait for his own or someone else's attempts to realize his ideas on an instrument or instruments. This is, though, no substitute for learning to play real instruments and to think in terms of them; rather it is a kind of "sounding scratch paper" where the student can find out "what happens if...". Indeed, such


experiments can bring him more quickly to an intelligent understanding of and ability to control and respond to pitch relationships, the interaction between pitch and duration, or the structure of melodies and more complex designs. This kind of understanding can, in turn, be quite directly transferred to learning to play a tune on a simple instrument, learning to read music and to models for composition or group improvisation. Computer implemented music should certainly not be a substitute for listening, playing and making music, live. Indeed, working with computer generated music seems to work rather as an incentive to make real maleable, human music.

The gadgetry is quite simple for the user: He finds himself before a typewriter which is coupled on one side with a computer (the coupling is by phone line, the computer, itself, is somewhere else) and on the other side with a "music box" about the size and shape of a lunch box. The music box can produce a five octave range of pitches and can play up to four parts simultaneously. The student "informs" the computer concerning the desired pitch, duration and some sort of operations which he wishes performed on his pitches and durations by using a typewriter or more accurately typewriter-terminal. The computer in turn causes the music box to produce the configurations indicated. The response is immediate. Thus in using the system the student need only be concerned with thinking out what processes he wants to happen. Paradoxically, he does not need to slip into the morass of becoming, himself, a music typewriter who types out notes on, let's say, the piano, without hearing how they go together or what they "mean". Using the electronic "music box" he can sit back and listen as often as he wishes to what he has invented, change it, and listen again to the result. Notice, especially, that while the student could be involved with getting a right answer (according to someone else's design) he is primarily concerned with

the results or effects of his own musical thinking. He might, for example, try to discover just what are the particular attributes of a melody that he finds "makes sense" to him and/or to his friends. Or he might want to find out how to make a "funny" melody, or how to turn a "straight" melody into a "silly" melody or a "scary" one. He might do any one of these, for example, by simply changing the set of durations he is working with, thereby transforming the character of some set of pitches that he has kept constant. The instant feedback of his ideas in sound and time tells him immediately the relation between musical means and effect: he has learned how to influence and control musical relationships, through designing a particular kind of musical process. A few specific examples may make the learning process more tangible.

On the basis of previous experience we have concluded that beginning students tend initially to hear whole configurations rather than discrete bits (like individual pitches); we thus prepared the following game as a starting point:

Using the typewriter terminal, students make the music-box play a complete familiar tune which we have programmed in advance--for example, "Twinkle Twinkle Little Star"--by simply typing STAR. The tune has also been previously broken down for him into the three phrases from which the whole melody can be built. Phrases we take, initially, to be the shortest perceptually accessible elements of a tune--analogous to particularly shaped building blocks which one can use to build a whole building. The student types B1 (which stands for Block 1 of the tune) and hears:



B2 gives him



B3 gives him:



The game is then to construct the whole melody out of these three basic building blocks entirely by listening to them in whatever order he chooses. The player can experiment with various orderings of the blocks: chains of blocks of any number and in any order can be "requested"--e.g., B1 B2 B3 B2 can be typed and immediately heard one after the other according to his choice and in time. Eventually, the student discovers how to build the tune out of its germinal tune blocks: B1 B2 B3 B3 B1 B2. Individuals of all ages seem to be captivated by this game but its purpose is more than simply to find the "solution". Consider what has happened:

- 1) The player is immediately involved in an active process, in listening and doing; at the same time, he is thinking of a melody too as an active process--indeed, one that can be built and described as a particular kind of procedure or active structure.
- 2) "Elements" of a tune are presented, initially, as perceptually accessible groupings (or phrases) derived from within the context of the tune, itself, and thus meaningful as structural events. This is in contrast to "elements" conceived as discrete events such as a pitch and its duration.
- 3) As a result of 2) the students are involved in an aural discrimination exercise that is context oriented rather than "absolute". That is, any pitch and its duration remains embedded in a grouping from which each individual event gains contextual "meaning".

- 4) Aural discrimination is thus comparatively general since it does not focus on individual events as if they were discrete entities. Students are asked to compare by ear the general "shape" of one tune block with the general shape of another tune block. For example, students recognize that BLOCK2 "sounds like an ending", or that BLOCK3 "has the same downward movement as BLOCK2 but it doesn't sound ended".
- 5) As the students experiment with various orderings, they discover that each ordering generates a different effect; indeed, that each tune block has a different effect or function depending on what comes before and after it. For example, BLOCK1 + BLOCK3 gives both blocks a new "meaning" as a result of the structural context. BLOCK3 + BLOCK2 reveals the parallel structure of these two blocks--something that might not have been noticed in the context of the original ordering.
- 6) Observations growing out of these first experiments lead students to ask a number of questions that are significant because they are so beautifully open-ended, so generously expandable. For example, "Why didn't I notice the parallel structure of B2 and B3 until I heard them in the reverse order while I was fooling around?" Or, "Why does the ordering B1 B3 B2, even though it seems ended, still sound incomplete, not self contained?" Or more simply, "Why does only B2 sound like an ending?" Any one of these questions would be difficult for a student to answer adequately at this initial phase in his musical development because all of them plunge the student into the intricacies of tonality as a system of interrelated


functions. But they are questions that the student can return to as his musical experiments continue and expand; when, later, for example, he can get into the "contents" of the tune blocks and can build strings of pitches and corresponding strings of durations, manipulating them in much the same way as he has manipulated the tune blocks. But most important, these are questions the student can find answers to himself by actually making things happen and watching (or really listening) to the results.

- 7) Almost as a by-product, the player of the game has, of course, found the structure or form of this melody as it would be described in conventional texts: A B A. But, to return to our first point, he has done so by building it as a process or procedurally -- he found out that the first two phrases return after a contrasting middle section. He also found out that unlike the two A sections, the B section creates contrast in part by repeating the same phrase twice, and he looked for that satisfying return after generating the incomplete middle section. In short, he has actively analyzed this melody!
- 8) Finally, the student's analysis should make it much easier for him to learn to play this or other tunes on a real instrument--say a recorder. He learns to play the three tune blocks already hearing, now, their similarities and differences, then simply follows the procedure he has discovered for ordering the tune blocks, and he knows the whole piece--in terms of its significant structural events--the "phrasing" is already part of his performance!


With this game as a beginning, students, themselves, think up other possibilities. We gave them unfamiliar tunes and later smaller segments--motives, rather than phrases. Students wanted to find tune blocks when

given a whole tune--which means, essentially, analyzing the phrase structure of the melody. This led to the possibility of students making up their own tune block games to be played by their friends. An interesting variation was that of providing students with just the tune blocks without giving them, first, the whole tune intact. With this game, the process is one of looking for an arrangement that "makes sense". The student's arrangement or procedure can be compared with the original composition, but what is really important here is, what does it mean to "make sense"? Again this is a question that can be explored on many levels--more of this in a moment. With the most advanced students, we used tunes where the segments were very similar to one another requiring careful discrimination among them; for example:

Beethoven Ländler



Perhaps more interesting is the possibility of constructing procedures for transforming motives or even a single motive and thereby building a whole melody. The most obvious of such procedures is sequential development. For example the following Polish folk song can be described entirely in terms of its initial motive. The procedure would take the first measure as a given and continue with a set of instructions for moving the initial motive up or down the appropriate interval:



The procedure would look something like this:

1. Play M (where M is the name for the germinal motive.)
2. Play M up 1
3. Play M down 1
4. Play M
- etc.

Or a canon can be written as a procedure. The principle of a three part piece in which all parts are identical can come alive: Given the "core" melody (say, "Row Row Row Your Boat") the student simply instructs the computer to do the following:

1. Make Voice 1: "ROW" (where "ROW" is the name for the whole tune)
2. Make Voice 2: "Rest 4 beats and then ROW"
3. Make Voice 3: "Rest 8 beats and then ROW"
4. Play Voice 1, Voice 2 and Voice 3 together.

While the students certainly have a good time making all this happen, notice that in doing so the student is discovering, through analysis, fundamental aspects of musical structure. The procedure for the Polish folk song is different from the round, which is different from the ABA, Twinkle Twinkle. But the analytic process is never passive, never merely visual (students work entirely by ear), never one of fitting some piece into a pre-determined mold. A design for building or transforming or combining melodies can be realized--that is, made to happen--as soon as the student sufficiently understands--that is, hears--the structural relations of the piece and can describe it procedurally. Hearing, Idea and action are always intertwined!

So far I have been describing activities where the "givens" are whole configurations--a phrase, a motive, a whole melody--which can be manipulated in various ways. What kinds of procedures might be involved in actually making a melody or in transforming the shape of a given phrase? A few examples from our work with both young children (fifth graders) and college students will give some idea of the possibilities. We started with just rhythmic configurations since these are more immediately accessible to beginning students. For this purpose the music box includes two non-pitch percussion sounds. Students began by playing the rhythm of a familiar tune, "Lightly Row", on a real drum.

! ! ! ! ! ! ! !

Asked to "draw a picture of the rhythm", most non-music-reading students did something like this:

Ex. 1: ||| ||| |||||


How about a duet with one person playing the "piece" and another "keeping time"--i.e., playing the beat that the piece generates. The picture of the duet (after experimenting with alternatives) looked like this:

Ex. 2: ||| ||| |||||
 | | | | | |

Now think of the beat as a constant--a measurer of time; the students assigned a number (12) to this constant, a "counter" with which to then measure the varied durations of the piece:

Ex. 3: 6 6 12 6 6 12 6 6 6 6 6 6 12
 \ \ \ \ \ \ \ \ \ \ \ \ \
 | | | | | | | |
 12 12 12 12 12 12 12 12


Of course the number assigned to the counter is more or less arbitrary; however, larger numbers indicate a slower beat (analogous to a larger distance), smaller numbers indicate a faster beat. The durations of the piece will, of course, maintain the same relationship to each other as well as to the beat when the basic counter changes. For example:



In this fashion students have gone from listening and playing, to various descriptions of what they have heard and played--a visual-spatial analogue of temporal relations (Examples 1 and 2) and a description of these same relations translated into numerical relations (Examples 3 and 4). So what started out as a kinesthetic-aural experience of a particular configuration is now concrete in another way: The rhythmic figure can be repeated, performed by others, it can be changed, added to, embedded in other sound environments, combined with pitches, etc. In addition the configuration has been analyzed in a way that is "extendable"--that is, the same kind of procedure can be used in dealing with other configurations so that the student can compare them, discover differences in character and structure and eventually learn to create the kind of figure he wants. For example: Invent a configuration that obscures a sense of beat; invent a figure that makes the beat appear, disappear and re-appear again; make a figure that causes the beat to group in twos, in threes, to shift from one to the other, etc.

In addition, the student can use his understanding to describe rhythmic figures in standard rhythm notation. The principle: All contiguous hits that together equal the constant time unit (the beat) form a group and are thus "hung together"

by a beam. The students' original drawings now translate into:





Once the student has discovered the possibility for describing a rhythm in terms of numerical relations, he can test his description by "asking" the music box to perform it. He can type:


Drum "6 6 12 6 6 12 6 6 6 6 6 6 12"

This will cause the electronic drum to play just what the student has been playing on his real drum! This done, the students can begin to explore the effect of a rhythmic figure when pitch is added. Initially, we gave the students the pitches of the melody, "Lightly Row", by pre-programming them and applying a name to the string of pitches. Thus, a student can type:


Sing :Row "6 6 12 6 6 12 6 6 6 6 6 6 12"

":Row" is the name for the string of pitches in the first long phrase of "Lightly Row". The numbers which follow are the durations which the students had already figured out and heard as a series of drum sounds. The new command, SING, causes the music box to play the first phrase of "Lightly Row" with each of its pitches given the proper duration. Again the "test" works-- the analysis now generates the tune!

Having thus reconstructed the tune, the students decided to invent an accompaniment figure for the tune to be played on the "drum". They tried two: "12 6 6" and "12 6 6 12" or  and .

The program they wrote said essentially: Drum "12 6 6" and keep on repeating it. The result was: .....

Listening to it played back, they were surprised; it sounded like:



This raised fundamental questions of grouping: What causes a set of durations with no change in loudness to "cluster" in some way? What generates an accent? These questions led to notions of meter, downbeats, and most important--what tends to generate duple grouping or triple grouping--i.e., duple or triple meter?


The questions became more alive and relevant when the students tried the second accompaniment figure (1 2 1 2 1 2 1 2). Some students heard the grouping as: 1 2 1 2 1 2 1 2 , others as: 1 2 1 2 1 2 1 2 but all agreed that the figure generated a 3 beat group in contrast to the two beat grouping of the first figure.

But the most dramatic effect occurred when each of the two accompaniment figures was played together with "Lightly Row". The first figure (1 2 1 2 1 2 1 2) played as an accompaniment made a piece that was "O.K., but not very lively or interesting"; the second accompaniment resulted in a piece that was "more varied, peppier, fun." What generates the difference in effect? This was not an easy question but some people (the more advanced students) discovered that the differences related to the fact that the 3 beat figure (1 2 1 2 1 2 1 2) was "out of phase" with the tune --i.e., the downbeats or accents did not coincide with the downbeats of "Lightly" nor did the accompaniment "come out right" with the ends of phrases. The higher level groupings were in conflict--duple against triple meter:

1 2 1 2 1 2 1 2
1 2 1 2 1 2 1 2

Thus, the students had at their command one possibility for creating differences in musical "character" or mood. Or putting it another way, by experimenting with possibilities, by making things happen and then exploring and questioning the results, students had learned something about the relation between musical means and effect. They had also discovered some rather fundamental aspects of musical structure and comprehensibility.

Jumping ahead a bit, now, the students later tried transforming the tune itself. With pitch and duration handled as interrelated but separable configurations, it was easy to keep the pitches constant and change just the durations:



Both the pitch configuration and the rhythmic pattern were made marvelously unrecognizable. Why? The new rhythmic figure caused a re-grouping of the pitches thus generating a different pitch-shape: Accents occur on different pitches than in the original version, the opening motive is broken up, etc. But isn't this just the kind of transformation that characterizes development or elaboration of a theme? Why not use this initial transformation to create a developmental contrasting section for "Lightly"? Further manipulation of pitch and duration finally led to the following piece:



We had a statement, a development based on the initial transformation of the pitch-duration material creating an effect of conflict or tension, followed by further elaboration and then return. I was reminded of the following passage from the Beethoven Piano Sonata, Op. 90, which occurs just before the recapitulation in the first movement:



But these examples are only a bare beginning. I have said little, for instance, about experimentation with pitch relations. It should be clear, however, that pitch relations can be made, transformed and manipulated in much the same way as rhythmic configurations. Our experience suggests, though, that it is crucial to give students the opportunity to deal with all 12 pitches from the beginning. They are then free to derive major scales, and to construct tonal functions from the all-pitch set if they wish, or to discover other bases for order and coherence. But it should be emphasized that the materials and means made available to the students allow them to deal with pitch and time as active, interrelated configurations. The interrelationship becomes a process to be heard and made into whole structures--whatever that may come to mean.

One more thing should be emphasized: The kind of activity I have described is only one part of musically productive learning. Clearly a

full music program must include listening to all kinds of music and learning to play an instrument. The whole point of the endeavor is lost if it does not lead to active participation, increased understanding and pleasure with great works as well as to an ability to make informed and appropriate ("musical") decisions as a performer. There must be continuous interaction between all these kinds of activity--not just from one month or year to the next, but every day. Speculation and theory developed at one's desk or together with the computer terminal may have a certain beauty but it must move out into the world and be practiced--in every sense of the word--before this kind of learning can become an integral part of one's life and breath. We are simply trying, here, to make learning a part of the world where music is really made; trying, that is, to erase the sometimes painful distinction between learning about music and doing it.